

(19)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 828 214 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
11.03.1998 Bulletin 1998/11

(51) Int Cl. 6: G06F 9/46, H04L 29/06

(21) Application number: 97306778.8

(22) Date of filing: 02.09.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE
Designated Extension States:
AL LT LV RO SI

(30) Priority: 09.09.1996 US 711189

(71) Applicant: LUCENT TECHNOLOGIES INC.
Murray Hill, New Jersey 07974-0636 (US)


(72) Inventor: Nepustil, Vladimir
Boulder, Colorado 80303 (US)

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)

(54) Dynamic reconfiguration of network servers

(57) One or more portions (221) of a database (220) which a primary server (106) uses to process client requests are duplicated (211) on one or more supplemental servers (105). As the clients demand for service increases and the processing load on the primary server becomes excessive (400), the primary server automatically off-loads the processing of those portions of the client requests that require one or more of the duplicated portions onto the supplemental servers by substituting (402) a secondary page (253) or a secondary object in

its database that points to the one or more duplicated portions in the supplemental servers for a corresponding primary page (252) or a primary object in its database that points to the one or more duplicated portions in its database. The supplemental servers then serve the portions of the client requests that require the one or more of the duplicated portions. As demand for service decreases and the primary server becomes under-loaded (404), it automatically restores (406) the primary page or the primary object in its database and resumes serving the entire client requests.



Description**Technical Field**

This invention relates to information network architecture, such as the Internet or an intranet architecture.

Background of the Invention

In an information network, such as the Internet, user's computers, referred to as clients, request information from information-providers' computers, referred to as servers, and the servers supply the requested information to the clients. In the World Wide Web (WWW), which is a de-facto standard for storing, finding, and transferring information on the Internet, the information is supplied in the form of pages. A page is a display screen-full of information expressed in textual, graphical, scriptural, and/or other form. A page comprises one or more information objects. An object is an information element that has its own network address -- preferably a unique single address-- called a URL (Uniform Resource Locator). For example, a page may comprise one or more text objects, one or more picture objects, and one or more script objects that are presented on the display screen in a layout defined by a frame object.

Typically, a server has a main page that serves as the entry point to the information and services that the server provides. This page typically points to other pages and to objects (e.g., graphic images, video/audio/text files, etc.), which are typically served by the same server.

Generally, when a client accesses the server, the server provides the main page to the client and then interacts with the client to provide the client with desired additional information and/or services. As increasing numbers of clients access the server, the server's processing load increases and its performance eventually degrades, so that users experience increasing delays between the time at which they place a request to the server and the time at which their request is satisfied by the server.

To avoid overloading of a server, typically an administrator must manually reconfigure the server and redirect some of the requests to other servers in order to lessen the load on the subject server. Some service providers store replicas of the served information in a plurality of servers and have different ones of the servers serve different requests, e.g., on a round-robin basis, thereby spreading the load of requests over multiple servers. This has several disadvantages. Firstly, an administrator's manual intervention is slow, inefficient, prone to error, and often not prompt. Secondly, using a plurality of servers to serve requests on a round-robin basis results in underutilization of the servers during periods when relatively few requests are being made, and hence it is inefficient. Furthermore, it requires all server information to be replicated on each server; the servers

cannot take advantage of a common cache for common data.

Summary of the Invention

This invention is directed to solving these and other problems and disadvantages of the prior art. Generally according to the invention, a portion of the information which a primary server uses to process client requests is replicated on one or more supplemental, stand-by, servers, and as the clients' demand for service increases and the processing load on the primary server becomes excessive, the primary server automatically off-loads the processing of those portions of the client requests that require the replicated portion of the information onto the supplemental servers. As demand for service decreases and the primary server becomes under-loaded, preferably the primary server automatically resumes serving the entire client requests.

The advantages of the invention include the following: the load-shedding and load-sharing happen automatically, without human intervention, based on the present processing load. Only one server, or one group of servers, out of the entire server set serves an individual portion of the information (e.g., a page, or an object, or a group of pages or objects) at any one time, which allows for efficient caching of the information. And more uniform response times are provided to clients even as client demand for services varies greatly. Moreover, while a standby server is not serving the primary server's clients, its processing power may be used for other processing activities, such as serving other clients whose demand for service peaks at a time different from the primary server's clients, thereby resulting in efficient server utilization.

According to a first aspect of the invention, a client-server system comprises a plurality of servers for processing client requests, wherein at least one first server of the plurality of servers has first information and second information related to the first information, for processing portions of the client requests that require the first information and portions of the client requests that require the second information. The at least one first server processes both portions of the client requests while the processing load on the at least one server is not excessive, e.g., does not exceed a predetermined first limit. In response to the processing load on the at least one first server becoming excessive, the at least one first server processes the portions of the client requests which require the first information without also processing the portions of the client requests which require the second information, and automatically redirects the portions of the client requests which require the second information to at least one second server for processing. The at least one server of the plurality of servers has the second information and processes the redirected portions of the client requests which require the second information, automatically in response to the



Europäisches Patentamt

European Patent Office

Office européen des brevets

(11)

EP 0 828 214 A3



(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
02.12.1998 Bulletin 1998/49

(51) Int Cl. 6: G06F 9/46, H04L 29/06

(43) Date of publication A2:
11.03.1998 Bulletin 1998/11

(21) Application number: 97306778.8

(22) Date of filing: 02.09.1997

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC
NL PT SE

(72) Inventor: Nepustil, Vladimir
Boulder, Colorado 80303 (US)

Designated Extension States:
AL LT LV RO SI

(74) Representative:
Watts, Christopher Malcolm Kelway, Dr. et al
Lucent Technologies (UK) Ltd,
5 Mornington Road
Woodford Green Essex, IG8 0TU (GB)


(30) Priority: 09.09.1996 US 711189

(71) Applicant: LUCENT TECHNOLOGIES INC.
Murray Hill, New Jersey 07974-0636 (US)

(54) Dynamic reconfiguration of network servers

(57) One or more portions (221) of a database (220) which a primary server (106) uses to process client requests are duplicated (211) on one or more supplemental servers (105). As the clients demand for service increases and the processing load on the primary server becomes excessive (400), the primary server automatically off-loads the processing of those portions of the client requests that require one or more of the duplicated portions onto the supplemental servers by substituting (402) a secondary page (253) or a secondary object in

its database that points to the one or more duplicated portions in the supplemental servers for a corresponding primary page (252) or a primary object in its database that points to the one or more duplicated portions in its database. The supplemental servers then serve the portions of the client requests that require the one or more of the duplicated portions. As demand for service decreases and the primary server becomes under-loaded (404), it automatically restores (406) the primary page or the primary object in its database and resumes serving the entire client requests.





European Patent
Office

EUROPEAN SEARCH REPORT

Application Number

EP 97 30 6778

| DOCUMENTS CONSIDERED TO BE RELEVANT | | | |
|--|--|-------------------|--|
| Category | Citation of document with indication, where appropriate, of relevant passages | Relevant to claim | CLASSIFICATION OF THE APPLICATION (Int.Cl.6) |
| A | ADLER R M: "DISTRIBUTED COORDINATION MODELS FOR CLIENT/SERVER COMPUTING" COMPUTER, vol. 28, no. 4, 1 April 1995, pages 14-22, XP000507856 * page 16, left-hand column, line 5 – page 18, left-hand column, line 30 * | 1-8 | G06F9/46 H04L29/06 |
| A | US 5 548 724 A (AKIZAWA MITSURU ET AL) 20 August 1996 * abstract * | 1,4,7,8 | |
| A | EP 0 648 038 A (IBM) 12 April 1995 * the whole document * | 1,4,7,8 | |
| TECHNICAL FIELDS SEARCHED (Int.Cl.6) | | | |
| G06F | | | |
| The present search report has been drawn up for all claims | | | |
| Place of search | Date of completion of the search | Examiner | |
| THE HAGUE | 12 October 1998 | Michel, T | |
| CATEGORY OF CITED DOCUMENTS | | | |
| X : particularly relevant if taken alone | T : theory or principle underlying the invention | | |
| Y : particularly relevant if combined with another document of the same category | E : earlier patent document, but published on, or after the filing date | | |
| A : technological background | D : document cited in the application | | |
| O : non-written disclosure | L : document cited for other reasons | | |
| P : intermediate document | & : member of the same patent family, corresponding document | | |

redirection. Preferably, the at least one first server automatically ceases redirecting the portions of the client requests that require the second information and resumes processing of both portions of the client requests in response to the processing load on the at least one first server ceasing to be excessive, e.g., falling below a predetermined second limit.

According to a second aspect of the invention, a method of operating a client-server system that includes a plurality of servers for processing client requests comprises the following steps. While a processing load on at least one first server of the plurality of servers is not excessive, the at least one first server processes both portions of client requests that require first information and portions of the client requests that require second information related to the first information; the at least one first server has both the first information and the second information. In response to the processing load on the at least one first server becoming excessive, the at least one server processes the portions of the client requests that require the first information without also processing the portions of the client requests that require the second information, and automatically redirects the portions of the client requests that require the second information to at least one second server of the plurality of servers. In response to the redirection, the at least one second server automatically processes the redirected portions of the client requests that require the second information; the at least one second server has the second information. Preferably, when the processing load on the at least one server falls below a predetermined limit, the at least one server automatically ceases to redirect the portions of the client requests that require the second information and resumes processing both portions of the client requests.

These and other advantages and features of the invention will become more apparent from the following description of an illustrative embodiment of the invention taken together with the drawing.

Brief Description of the Drawing

FIG. 1 is a block diagram of an information network that includes an illustrative embodiment of the invention;

FIG. 2 is a block diagram of partial contents of memories of servers of the information network of FIG. 1; and

FIGS. 3-5 each are a flow diagram of partial operations of a different one of the servers of the information network of FIG. 1.

Detailed Description

FIG. 1 shows an illustrative information network, which in this example is the Internet. It comprises a plurality of clients 100-101 and a plurality of servers 105-107 interconnected by the Internet network fabric

110. Each server 105 is a computer having a processor and a memory, wherein the processor executes control programs stored in the memory to provide services and data stored in the memory. Each server 105-107 is a primary server for a database of information A-C, respectively. According to the invention, however, besides being a primary server for a database, each server 105-107 is also a secondary, or supporting, server for a portion of one or more other servers' databases. In the example of FIG. 1, server 105 is a supplemental server for a portion BZ of server's 106 database B; server 106 is a supplemental server for a portion AY of server's 105 database A and a portion CW of server's 107 database C; and server 107 is a supplemental server for a portion AX of server's 105 database A.

FIG. 2 shows the data contents of memories 205-207 of servers 105-107, respectively, that are relevant to an understanding of this invention. Memory 205 stores database A 210 and a duplicate 211 of portion BZ 221 of database B 220. Memory 206 stores database B 220, a duplicate 222 of portion CW 231 of database C 230; and a duplicate 223 of portion AY 213 of database A 210. And memory 207 stores database C 230 and a duplicate 232 of portion AY 212 of database A 210.

Database A 210 comprises a primary main page A1 240 that includes links to other pages A2-A4 243-245 in database A 210. Pages A3 244 and A4 245 in turn include links to yet other pages A5 246 and A6-A7 247-248, respectively, in database A 210. Database A 210 further includes a secondary main page A1' 241 that includes links to pages A2-A3 243-244 in database 210, but instead of including a link to page A4 245 of portion AX 212 in memory 205, it includes a link to duplicate page A4 245 of duplicate portion AX 232 in memory 207. Database A 210 yet further includes a tertiary main page A1" 242 that includes a link to page A2 243 in database 210, but instead of including links to pages A3 244 and A4 245 of portions AX 212 and AY 213 in memory 205, it includes links to duplicate pages A3 244 and A4 245 of duplicate portions AY 223 and AX 232, respectively, in memories 206 and 207, respectively.

Database B 220 comprises a main page B1 250 that includes links to an object B2 251 and to another page B3 252 in database B 220. B3 252 is a primary page that in turn includes links to other objects and/or pages B4-B5 258-259 in database B 220. Database B 220 further includes a secondary page B3' 253 which, instead of including links to pages B4-B5 258-259 in portion BZ 221 of database B 220 in memory 206, includes links to duplicate objects and/or pages B4-B5 258-259 of duplicate portion BZ 211 in memory 205.

Database C 230 comprises a primary format-object C1 260 of a main page that includes a link to a data object C2 262 in portion CW 231 of database C 230. Database C 230 further includes a secondary format-object C1' 261 of the main page which, instead of including a link to data object C2 262 in portion CW 231 of database C 230 in memory 207, includes a link to duplicate object

C2 262 of duplicate portion CW 222 in memory 206.

It is assumed that servers 105-107 experience their heaviest processing loads at different times. For example, server 105 may be most heavily used in the evenings, server 106 may be most heavily used on workdays, and server 107 may be most heavily used on weekends. Initially, all servers 105-107 operate conventionally. That is, server 105 initially serves all requests for information from database A 210, server 106 initially serves all requests for information from database B 220, and server 107 initially serves all requests for information from database C 230, in a conventional manner, using primary pages and/or objects 240, 252, and 260. As is also conventional, each server 105-107 keeps a record of its present processing load, for example, in the form of a number of accesses (requests) served per unit of time.

Operations of servers 105-107 that are relevant to an understanding of the invention are diagrammed in FIGS. 3-5, respectively. In addition to its conventional operations, each server 105-107 executes a load-control program that is initialized with predetermined load limits. As shown in FIG. 3, server 105 repeatedly checks if its present processing load exceeds a high-load upper limit "A1' high" at step 300. If not, it means that server 105 is not overloaded, and so server 105 remains at step 300; if so, it means that server 105 is overloaded, and server 105 therefore substitutes secondary page A1' 241 for primary page A1 240, at step 302. This has the effect of causing all subsequent requests for information from portion AX 212 of database A 210 to be directed to server 107. Server 107 serves these requests in a conventional manner from duplicate portion AX 232. Hence, when server 105 becomes overloaded, some of its processing load is taken over by server 107.

Following step 302, server 105 checks if a high-load lower limit "A1' low" exceeds its present processing load, at step 304. If so, it means that server 105 is underloaded, and so server 105 substitutes primary page A1 240 for secondary page A1' 241, at step 306. This has the effect of resuming initial operation, where server 105 is serving all requests for information from database A 210. Server 105 then returns to step 300.

If it is determined at step 304 that limit A1' low does not exceed the present load, server 105 checks whether the present load again exceeds the A1' high limit, at step 308. If not, it means that server 105 is not overloaded, and so server 105 returns to step 304; if so, it means that server 105 is again overloaded, and server 105 therefore substitutes tertiary page A1' 242 for secondary page A1' 241, at step 310. This has the additional effect of causing all subsequent requests for information from portion AY 213 of database A to be directed to server 106. Server 106 serves these requests in a conventional manner from duplicate portion AY 223, thereby taking on some of the processing load that would otherwise have to be done by server 105, and hence reducing the load on server 105.

Following step 310, server 105 checks if the A1' low limit exceeds the present processing load of server 105, at step 312. If not, server 105 remains at step 312; if so, it means that server 105 is underloaded, and server 105 therefore substitutes secondary page A1' 241 for tertiary page A1' 242, at step 314. This has the effect of server 105 taking back the portion of the processing load that had been transferred to server 106 at step 310. Server 105 then returns to step 304.

The operation of servers 106 and 107 is similar. As shown in FIG. 4, server 106 repeatedly checks if its present processing load exceeds a high-load upper limit "B3' high", at step 400. If not, server 106 is not overloaded and remains at step 400; if so, server 106 is overloaded, and therefore it substitutes secondary page B3' 253 for primary page B3 252, at step 402. This has the effect of causing all subsequent requests for information from portion BZ 221 of database B 220 to be directed to server 105. Server 105 serves these requests in a conventional manner from duplicate portion BZ 211, thereby relieving the load on server 106.

Following step 402, server 106 checks if a high-load lower limit "B3' low" exceeds its present processing load at step 404. If not, server 106 remains at step 404; if so, it means that server 106 is underloaded, and therefore server 106 substitutes primary page B3 252 for secondary page B3' 253, at step 406. This has the effect of resuming initial operation, where server 106 is serving all requests from database B220. Server 106 then returns to step 400.

As shown in FIG. 5, server 107 repeatedly checks if its present processing load exceeds a load limit "C1'", at step 500. If not, server 107 is not overloaded and remains at step 500; if so, server 107 is overloaded, and therefore it substitutes secondary object C1' 261 for primary object C1 260, at step 502, thereby transferring some of its processing load to server 106. Following step 502, server 107 repeatedly checks if the load limit "C1'" exceeds its present processing load, at step 504. If not, server 107 remains at step 504; if so, it means that server 107 is no longer overloaded, and therefore server 107 substitutes primary object C1 260 for secondary object C1' 261, at step 506, thereby resuming its initial operation. Server 107 then returns to step 500.

Of course, various changes and modifications to the illustrative embodiment described above will be apparent to those skilled in the art. For example, instead of storing both primary and secondary pages or objects, the primary pages or objects can be converted into the secondary pages or objects "on-the-fly" (e.g., in real time), and vice versa. Likewise, instead of duplicate portions of the database being pre-stored on supplemental servers, the database portions may be duplicated and distributed to the supplemental servers "on-the-fly." Moreover, measurements and limits other than the number of accesses per unit of time can be used to determine whether to off-load or return processing from or to the primary server. These measurements and limits

can be forward-looking, such as predictive algorithms which estimate future load based on load experienced at a similar time in the past. Furthermore, the main server can request present processing load data from the stand-by servers and incorporate these data into its decision of whether to off-load processing to those stand-by servers.

Claims

5

10

15

20

25

30

35

40

45

50

55

sponse to the processing load on the at least one first server becoming excessive.

4. The system of claim 2 wherein:

the at least one first server (105) processes both portions of the client requests while the processing load on the at least one first server does not exceed (500) a predetermined limit, and processes the portions of the client requests which require the first information without also processing the portions of the client requests which require the second information and automatically redirects (502) the portions of the client requests which require the second information in response to the processing load on the at least one first server exceeding (500) the predetermined limit.

5. The system of claim 4 wherein:

said predetermined limit is a first predetermined first limit, and the at least one first server automatically ceases (506) redirecting the portions of the client requests that require the second information and resumes processing of both portions of the client requests in response to the processing load on the at least one first server falling below (504) a predetermined second limit.

6. The system of claim 1 wherein:

the at least one first server (105) further has third information (213) related to the first information (210), processes the portions of the client requests that require the first information (210), the portions of the client requests that require the second information (212), and portions of the client requests that require the third information (213) while the processing load on the at least one server does not exceed (300) a predetermined first limit, processes the portions of the client requests which require the first information and the portions of the client requests which require the third information without also processing the portions of the client requests which require the second information and automatically redirects (302) the portions of the client requests which require the second information to the at least one second server for processing in response to the processing load on the at least one first server exceeding (300) the predetermined first limit for a first time, and processes the portions of the client requests which require the first information without also processing the portions of the client requests which require the second information and the portions of the client requests which require the third information and automatically redirects (310) the portions of the client requests which require the third information

1. A client-server system (FIG. 1) comprising:

a plurality of servers (105-107) for processing client (100-101) requests;
at least one first server (105) of the plurality of servers having first information (210) and second information (212) related to the first information, for processing portions of the client requests that require the first information and portions of the client requests that require the second information, while a processing load on the at least one server is not excessive, CHARACTERISED BY the at least one first server processing portions of the client requests which require the first information (210) without also processing the portions of the client requests which require the second information (212) and automatically redirecting the portions of the client requests which require the second information to at least one second server (107) for processing, in response to the processing load on the at least one first server becoming excessive; and

the at least one second server (107) of the plurality of servers having the second information (232) of the first and the second information, for processing the redirected portions of the client requests which require the second information, automatically in response to the redirection.

2. The system of claim 1 wherein:

the at least one second server (107) has the second information (232) without also having the first information (210), for processing the redirected portions of the client requests which require the second information without also processing the portions of the client requests which require the first information.

3. The system of claim 1 wherein:

the first information includes a primary page or object (240) that points to the second information (212) of the at least one first server (105) while the processing load on the at least one first server is not excessive, and includes a secondary page or object (241) that points to the second information (232) of the at least one second server (107) in re-

FIG. 2





FIG. 3



THIS PAGE BLANK (USPTO)